

UNITED STATES PATENT APPLICATION

for

COMMUNICATING MESSAGE REQUEST TRANSACTION
TYPES BETWEEN AGENTS IN A COMPUTER SYSTEM
USING MULTIPLE MESSAGE GROUPS

Inventor:

David Harriman

prepared by:

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP
12400 Wilshire Boulevard
Los Angeles, CA 90025-1026
(408) 720-8300

File No.: 42390.P13766

EXPRESS MAIL CERTIFICATE OF MAILING


"Express Mail" mailing label number: EL617207646US

Date of Deposit: December 28, 2001

I hereby certify that I am causing this paper or fee to be deposited with the United States Postal Service "Express Mail Post Office to Addressee" service on the date indicated above and that this paper or fee has been addressed to the Assistant Commissioner for Patents, Washington, D. C. 20231

MELANIE LYONS

(Typed or printed name of person mailing paper or fee)



(Signature of person mailing paper or fee)

12-28-01

(Date signed)

1004055-122801
T0822T-5520400T

**COMMUNICATING MESSAGE REQUEST TRANSACTION TYPES BETWEEN AGENTS IN A
COMPUTER SYSTEM USING MULTIPLE MESSAGE GROUPS**

Field Of The Invention

[0001] The present invention pertains to the field of computer systems. More particularly, this invention pertains to the field of high speed point-to-point interconnections and communications architectures.

Background of the Invention

[0002] Computing appliances, e.g., computer systems, servers, networking switches and routers, wireless communication devices, and the like are typically comprised of a number of disparate elements. Such elements often include a processor, system control logic, a memory system, input and output interfaces, and the like. To facilitate communication between such elements, computing appliances have long relied on general purpose input/output busses to enable these disparate elements of the computing system to communicate with one another in support of the myriad of applications offered by such appliances.

[0003] Perhaps one of the most pervasive of such general purpose bus architectures is the Peripheral Component Interconnect (PCI) bus. The PCI bus standard (Peripheral Component Interconnect (PCI) Local Bus Specification, Rev. 2.2, released December 18, 1998) defines a multi-drop, parallel bus architecture for interconnecting chips, expansion boards, and processor/memory subsystems in an arbitrated fashion within a computing appliance. While typical PCI bus implementations have a 133Mbps throughput (i.e., 32bits at 33MHz), the PCI 2.2 standard allows for 64bits per pin of the parallel connection clocked at up to 133MHz resulting in a theoretical throughput of just over 1Gbps.

[0004] The throughput provided by the PCI bus architectures has, until recently, provided adequate bandwidth to accommodate the internal communication needs of even the most advanced of computing appliances (e.g., multiprocessor server applications, network appliances, etc.). However, recent advances in processing power and increasing input/output bandwidth

demands create a situation where prior general purpose architectures such as the PCI bus architecture have become processing bottlenecks within such computing appliances.

[0005] Another limitation associated with prior architectures is that they are typically not well-suited to process isochronous (time dependent) data streams. An example of an isochronous data stream is a multimedia data stream which requires a transport mechanism to ensure that the data is consumed as fast as it is received and to ensure that the audio portion is synchronized with the video portion. Conventional general purpose input/output architectures process data asynchronously, or in random intervals as bandwidth permits. Such asynchronous processing of multimedia streams data can result in lost data and/or misaligned audio and video.

10040755-12301

Brief Description of the Drawings

[0006] The invention will be understood more fully from the detailed description given below and from the accompanying drawings of embodiments of the invention which, however, should not be taken to limit the invention to the specific embodiments described, but are for explanation and understanding only.

[0007] Figure 1 is a block diagram of one embodiment of a computer system.

[0008] Figure 2 is a graphical illustration of an example enhanced general input/output port.

[0009] Figure 3 is a diagram showing the format of one embodiment of the start of a transaction layer packet header.

[0010] Figure 4 is a diagram of a request packet header supporting a 32-bit address format.

[0011] Figure 5 is a diagram of a request packet header supporting a 64-bit address format.

[0012] Figure 6 is a diagram of a packet header for a Message.

[0013] Figure 7 is a diagram showing a request header format for a configuration transaction.

[0014] Figure 8 is a diagram showing one embodiment of a format for a completion header.

[0015] Figures 9a and 9b combined form a flow diagram of an example embodiment of a method for handling received transaction layer packets.

[0016] Figure 10 is a flow diagram of one embodiment of a method for handling error conditions associated with received request packets.

[0017] Figure 11 is a flow diagram of one embodiment of a method for handling a completion packet that is not expected by a system agent.

[0018] Figure 12 is a flow diagram of one embodiment of a method for a requesting device handling a completion packet with a completion status other than "Successful Completion."

[0019] Figure 13 is a flow diagram of one embodiment of a method for a completing device handling a completion packet with a completion status other than "Successful Completion."

4040355-12801

Detailed Description

[0020] Described below are embodiments of a point-to-point packet-based interconnection architecture, communication protocol and related methods to provide a scalable and extensible general input/output communication platform for deployment within an electronic appliance. The disclosed embodiments involve an enhanced general input/output interconnection architecture and associated communications protocol. One example embodiment includes one or more of a root complex including a host bridge, a switch, or endpoints, each incorporating at least a subset of enhanced general input/output features to support enhanced general input/output communication between such elements.

[0021] Communication between the enhanced general input/output facilities of such elements is performed in one embodiment using serial communication channels employing a communication protocol which supports one or more innovative features including, but not limited to, virtual communication channels, trailer-based error forwarding ("trailers" are appended to transaction layer packets to indicate an error condition), support for legacy PCI-based devices, multiple request response types, flow control and/or data integrity management facilities. The communication protocol supported in this embodiment includes a communication protocol stack including a physical layer, a data link layer and a transaction layer.

[0022] In an alternate embodiment, a communications agent incorporates an enhanced general input/output engine comprising a subset of the foregoing features. Further, one or more elements of the various embodiments may be implemented in hardware, software, a propagated signal, or a combination thereof.

[0023] Figure 1 is a block diagram of an electronic appliance 100, which, for this embodiment, is a computer system. The system 100 includes a processor 102, a host bridge 103 included as part of a root complex 104, switch 108 and endpoint 110, each coupled as shown. The root complex 104, switch 108, and endpoint 110 include one or more instances of an enhanced general input/output communication port 106. As shown, each of the elements 102,

104, 108 and 110 are coupled to at least one other element through a communication link 112 supporting one or more enhanced general input/output communication channels via the enhanced general input/output communication port. The system 100 is intended to represent any of a wide variety of traditional and non-traditional computing systems, servers, network switches, network routers, wireless communication subscriber units, wireless communication telephony infrastructure elements, personal digital assistants, set-top boxes, or any electric appliance that would benefit from the communication resources introduced through integration of at least a subset of the enhanced general input/output interconnection architecture and/or communications protocol described herein.

[0024] In this example embodiment, processor 102 controls one or more aspects of the functional capability of the electronic appliance 100. In this regard, the processor 102 is representative of any of a wide variety of control logic devices including, but not limited to, one or more of a microprocessor, a programmable logic device (PLD), programmable logic array (PLA), application specific integrated circuit (ASIC), a microcontroller, and the like.

[0025] The root complex 104 provides a communications interface between the processor 102 and the switch 108 and endpoint 110. As used herein, the term "root complex" refers to a logical entity of an enhanced general input/output hierarchy that is closest to a host controller, a memory controller hub, an IO controller hub, or any combination of the above, or some combination of chipset/CPU elements (i.e., in a computing system environment). Although depicted in Figure 1 as a single unit, the root complex 104 may be implemented with multiple physical components. The root complex 104 is populated with one or more enhanced general input/output ports 106 to facilitate communication with other peripheral devices, e.g., the switch 108, endpoint 110 and, although not particularly depicted, the legacy bridges 114, or 116. In one embodiment, each enhanced general input/output interface port represents a different hierarchy domain. In this regard, the embodiment of Figure 1 denotes a root complex 104 with three hierarchy domains.

[0026] Figure 2 is a graphical illustration of an example enhanced general input/output port 106. In this embodiment, the enhanced general input/output port 106 implements a communication stack comprising a transaction layer 202, a data link layer 204 and a physical layer 206 including a logical sub-block 208, and a physical sub-block 210, as shown. Elements of the transaction layer will be discussed below in more detail.

[0027] The transaction layer 202 provides an interface between the enhanced general input/output architecture and a device core. A primary responsibility of the transaction layer 202 is the assembly and disassembly of packets for one or more logical devices within an agent.

[0028] One of the primary goals of the enhanced general input/output architecture is to maximize the efficiency of communication between devices. In one embodiment, the transaction layer implements a pipelined full split-transaction protocol as well as mechanisms for differentiating the ordering and processing requirements of transaction layer packets. The transaction layer further comprehends transaction layer packet construction and processing.

[0029] One embodiment of the enhanced general input/output architecture supports the following basic transaction types and address spaces: Memory, I/O, Configuration, and Message. Two addressing types are supported: 32bit and 64bit.

[0030] Transactions are carried using Request and Completion packets, which may be referred to simply as Requests and Completions. Completions are used only where required, e.g.: to return read data, or to acknowledge completion of I/O and configuration write transactions. Completions are associated with their corresponding Requests by the value in the Requester ID field of the packet header (discussed below).

[0031] All transaction layer packets in this embodiment start with a defined header. Some transaction layer packets include data following the header as determined by the format field specified in the transaction layer packet header. The transaction layer packet is limited in size by a predetermined maximum payload size value. The transaction layer packet data in this embodiment is four-byte naturally aligned and in increments of four-byte double-words.

[0032] Figure 3 is a diagram showing the format of one embodiment of the start of a transaction layer packet header. Each transaction layer packet header includes a three-bit format field (Fmt[2:0]). The transaction layer packet header also includes a four-bit type field (Type[3:0]). Both the Fmt and Type fields need to be decoded in order to determine the transaction layer packet format. Table 1 below shows example encodings for the Fmt field.

000	2 double-word header, no data
001	3 double-word header, no data
010	4 double-word header, no data
101	3 double-word header, with data
110	4 double-word header, with data

Table 1 – Fmt Field Encodings

[0033] The transaction layer header for this embodiment also includes a two-bit Extended Type/Extended Length field (Et/El). This field is used to extend either the Type field or the Length field, depending on the value in the Type field. The Length field for this embodiment is ordinarily an eight-bit field, but may be extended to become a ten-bit field if the value in the Type field indicates that the Et/El field is to be used to extend the Length field. The Type field can be extended to become a six-bit field by appending the Et/El field, depending on the value in the Type[3:0] field. See Table 2 below for example Fmt, Type, and Et/El field encodings (alternative embodiments may use other encoding schemes). The Et/El field is used as an extension of the Type field except where noted.

Packet Type	Fmt[2:0]	Type[3:0]	Et/El[1:0]	Description
MRd	001 010	1001	El9 El8	Memory Read Request Et/El field used for Length[9:8]
MRdLk	001 010	1011	00	Memory Read Request - Locked
MWr	101 110	0001	El9 El8	Memory Write Request - Posted

				Et/El field used for Length[9:8]
IORd	001	1010	00	IO Read Request
IOWr	101	1010	00	IO Write Request
CfgRd0	001	1010	01	Configuration Read Type 0
CfgWr0	101	1010	01	Configuration Write Type 0
CfgRd1	001	1010	11	Configuration Read Type 1
CfgWr1	101	1010	11	Configuration Write Type 1
Msg	010	011s ₂	s ₁ s ₀	Message Request – The sub-field s[2:0] specifies a group of messages. The Message field must be decoded to determine specific cycle including if a Completion is required or not.
MsgD	110	011s ₂	s ₁ s ₀	Message Request with Data - The sub-field s[2:0] specifies a group of messages. The Message field must be decoded to determine specific cycle including if a Completion is required or not.
MsgComm	110	110c ₂	c ₁ c ₀	Message for Advanced Switching – The sub-field c[2:0] specifies the message type: 000 - Unicast, Data Packet 001 - Multicast, Data Packet 010 - Signaling Packet, without interrupt 011 - Reserved 100 - Null signaling Packet, interrupt to Host in the destination Hierarchy 101 - Null signaling Packet, interrupt to destination device 110 - Signaling Packet, with interrupt to Host in the destination Hierarchy 111 - Signaling Packet, with interrupt to destination device
Cpl	001	0100	00	Completion without Data – used for IO and Configuration Write Completions, and Memory Read Completions

				with Completion Status other than "Successful Completion."
CplD	101	0100	E19 E18	Completion with Data – used for Memory, IO, and Configuration Read Completions Et/EI field used for Length[9:8]
CplDLk	101	0101	01	Completion for Locked Memory Read

Table 2 – Fmt, Type, and Et/EI Encodings

[0034] Request packets include a Request Header which for some types of Request packets will be followed by some number of double-words of data. The term "double-word" as used herein indicates a 32-bit length of data. For this example embodiment, the Length field for Message Request Headers is not used except for Messages that explicitly refer to a data length. Also for this embodiment, for Memory Read Requests and Memory Write Requests, the EI/Et field is concatenated with the Length field to form a ten-bit length field. The ten-bit length field allows read and write requests indicating up to 4kB of data. Other types of transaction layer packets are limited by the size of the Length[7:0] field to indicating up to 1kB of data. The amount of data included in any transaction layer packet is limited in one embodiment to a predetermined maximum payload size. For transaction layer packets that include data, the value in the Length field and the actual amount of data should be equal. If the receiver determines that the Length field value and the actual amount of data do not match, then the packet is treated as a Malformed Transaction Layer Packet. Malformed Transaction Layer Packets are described below.

[0035] Figure 4 is a diagram of a request packet header supporting a 32-bit address format and Figure 5 is a diagram of a request packet header supporting a 64-bit address format. For one embodiment, memory read requests and memory write requests case use either the 32-bit address format or the 64-bit address format. For addresses below 4GB, the 32-bit format is used.

[0036] The request packet headers of Figures 4 and 5 also include a First Double Word Byte Enable field (1st DW BE) and a Last Double Word Byte Enable field (Last DW BE). The First Double Word Byte Enable field contains byte enables for the first double word of any memory read or write request. This field also contains byte enables for the only double word of an input/output or configuration request. The Last Double Word Byte Enable field contains byte enables for the last double word of any memory read or write request. The byte enable fields are not used with Messages because these fields overlap the message code field for a message request header (See Figure 7, discussed below).

[0037] For one embodiment, for each bit in the byte enable fields, a value of "0" indicates that the corresponding byte of data is not written or, if non-prefetchable, read at a completor. The term "completor" as used herein is meant to indicate a logical device addressed by a request packet header. A value of "1" indicates that the corresponding byte of data is written or, if non-prefetchable, read at the completor. For the First Double Word Byte Enables field, bit 0 corresponds to byte 0 of the first double word of data. Bit 1 corresponds to byte 1 of the first double word of data. Bit 2 corresponds to byte 2 of the first double word of data. Bit 3 corresponds to byte 3 of the first double word of data. For the Last Double Word Byte Enables field, bit 0 corresponds to byte 0 of the last double word of data. Bit 1 corresponds to byte 1 of the last double word of data. Bit 2 corresponds to byte 2 of the last double word of data. Bit 3 corresponds to byte 3 of the last double word of data.

[0038] The example packet headers of Figures 4, 5, 6, and 8 include a Requestor ID field, a Tag field, an Attribute field, and a Virtual Channel ID field. The Requestor ID field and the Tag field together form a Transaction ID field. The Requestor ID field is divided into a Bus Number field, a Device Number field, and a Function Number field.

[0039] The Tag field is a 5-bit field generated by each requesting device. The tag value is unique for all outstanding requests that require a completion for that requesting device. The Transaction ID field is included with all Requests and Completions. The Requestor ID field for these example embodiments is a 16-bit value that is unique for every function (a function is one

independent section of a multi-function device identified in configuration space by a unique function number). Functions capture the Bus Number supplied with all configuration writes completed by the function, and supply this number in the Bus Number section of the Requestor ID field. Each logical device in a component is designed to respond to a unique Device Number for configuration requests addressing that component. For these example embodiments a component may contain many (perhaps up to several dozen) logical devices. Each function associated with a logical device in a component is designed to respond to a unique function number for configuration requests addressing that component and logical device. Each logical device may contain up to eight logical functions.

[0040] The Attributes field specifies characteristics of the transaction. Attributes that may be specified in the attributes field include a priority attribute, transaction ordering attributes, and cache coherency management attributes.

[0041] The Virtual Channel ID field identifies the virtual channel. For these example embodiments, the Virtual Channel ID field is a 4-bit field that allows identification of up to 16 virtual channels on a per transaction basis. For these example embodiments, virtual channel 0 is used for general purpose traffic and a virtual channel other than 0 is used for isochronous traffic.

[0042] Figure 6 is a diagram of a packet header for a Message. As seen in Table 2, messages may or may not include data and may or may not require completion. Messages are decoded by all devices in a system that support the enhanced general input/output interconnection architecture.

[0043] For message requests, the Message field is decoded in order to determine the specific cycle and to determine whether the message includes data and whether the message requires completion. The Message field for this embodiment is an 8-bit field located where the byte enable fields normally reside for other transaction types. Unsupported messages are treated by the receiving device as no-completion-required (no-completion-required transactions discussed below).

[0044] Messages for this example embodiment are divided into groups. There are eight groups that include data with the request and eight groups that do not. Other embodiments are possible using different numbers of groups. For this embodiment, as shown in Table 2, the eight groups that include data with the requests have a value of b110 in the Fmt field. The eight groups that do not include data have a value of b010 in the Fmt field. The sub-field s[2:0] incorporates one bit from the Type field and the two bits from the Et/El field. The sub-field s[2:0] indicates one of eight groups.

[0045] Examples of various messages that may be implemented include, but are not limited to, the following: messages for unlocking devices; messages for resetting devices; messages indicating a correctable error condition; messages indicating an uncorrectable error condition; messages indicating a fatal error condition; messages used to report bad request packets; messages relating to power management; messages related to ordering control/management; and messages for emulating legacy (e.g., PCI) interrupt signals (or other legacy signals). These various message types can be subdivided into one of the previously discussed groups. For example, all of the power management messages may be included in one group and the interrupt signaling messages may be included in another. Embodiments are also possible where one or more groups are set aside for vendor-specific use.

[0046] Figure 7 is a diagram showing a request header format for a configuration transaction. The configuration space is one of the four supported address spaces for these example embodiments.

[0047] Figure 8 is a diagram showing one embodiment of a format for a completion header. All read requests and some write requests require completion. Completions include a completion header that, for some types of completions, are followed by some number of double-words of data. The Completion Status[2:0] field shown in Figure 8 indicates the status for a completion. Table 3 shows one example encoding scheme.

Completion Status[2:0]	Status
000	Successful Completion
001	Unsupported Request – Expecting Completion
011	Reserved
100	Completer Abort

Table 3 – Completion Status Field Encoding Scheme

[0048] The Completer ID[15:0] field contains the same type of information as the Requestor ID field described above. The value provided in the Completer ID field corresponds to the bus/device/function of the agent that completes the request. Completion headers contain the same values for Requester ID, Tag, and Channel ID as were supplied in the header of the request packet. Completion headers also contain the same value in the Attribute field as was initially supplied with the header of the request. Completion packets are routed by switches and root complexes to the port that initiated the corresponding request transaction.

[0049] For memory read request transactions, individual completion packets may provide less than the full amount of data requested by the corresponding read request so long as all of the completion packets associated with the corresponding read request, when combined, return the amount of data specified. For these example embodiments, I/O and Configuration read requests are completed with exactly one completion packet.

[0050] A completion that includes data specifies the amount of data in the packet header. If the completion packet actually contains an amount of data that differs from the specified amount, a malformed transaction layer packet results.

[0051] Figures 9a and 9b combined form a flow diagram of an example embodiment of a method for handling received transaction layer packets. The operations described below do not necessarily need to occur in a serial fashion. Some embodiments may perform some operations simultaneously. At block 905, a check is made to determine whether the values contained in the Fmt and Length fields of the received packet match the actual size of the packet. A mismatch indicates a malformed packet and an error case results as shown at block 925. Error case

handling will be discussed below. If the actual size of the received packet does not indicate a mismatch with the Fmt and Length fields, then processing continues at block 910.

[0052] If the received packet is a memory request using 64 bit addressing, then at block 910 the address bits [63:32] are checked to see whether any of the address bits [63:32] are non-zero. If none of the address bits [63:32] are non-zero, then the result is a malformed packet and processing proceeds to the error case block 925. If at least one of the address bits [63:32] are non-zero, then processing continues at block 915.

[0053] At block 915, a check is made to determine whether any fields in the packet header contain reserved values. If reserved values are found, then the result is an malformed packet and processing proceeds to block 925. If no reserved values are found, then processing continues at block 930.

[0054] At block 930, a determination is made as to whether the packet is a request packet or a completion packet. If the packet is a completion packet, then processing proceeds to the completion handling block 935. Completion handling will be discussed more fully below. If the received packet is not a completion packet, then processing continues at block 940. Note that all packets that flow to block 940 are request packets.

[0055] At block 940, a check is made to determine whether the request packet is a request type that is supported by the completing device. If the request type is not supported, then the result is an unsupported request and processing proceeds to the error case block 925. For this example embodiment, if the unsupported request type is a broadcast message or a message using an encoding reserved for broadcast messages, then the packet is silently dropped and no error case results. If the request type is supported by the completing device, then processing continues at block 945.

[0056] If, as shown at block 945, the completing device is unable to respond to the request packet due to an internal error, then the result is a "completer abort" and processing proceeds to error case block 925. Otherwise, the request is serviced at block 950. In servicing the request, it may be necessary to repeat the processing indicated by blocks 940 and 945.

[0057] Once the request is serviced successfully, then processing continues at block 955. As indicated by block 955, if the processed request requires a completion, then a completion packet is returned at block 960.

[0058] Figure 10 is a flow diagram of one embodiment of a method for handling error conditions associated with received request packets. As seen at block 1010, if the received request is expecting a completion, then a completion with the appropriate completion status is transmitted at block 1020. The completion is routed back to the requesting device. If the received request is not expecting a completion, then an error message is transmitted to the requesting device at block 1030. The error is reported to the system at block 1040. The error message transmitting operating denoted at block 1030 may be implemented as a programmable option.

[0059] Some systems may include one or more PCI buses in addition to the previously discussed enhanced general input/output interconnection architecture. For memory, I/O, and configuration requests traveling through the enhanced general input/output interconnection architecture and destined to a device on a PCI bus, the completion status for these embodiments represents the actual PCI termination for the cycle. For example, a non-posted PCI cycle must actually be serviced on the PCI bus before a completion status can be determined. For all other cases, the completion status values are defined as discussed below.

[0060] When a request is completed successfully by the completing device, the resulting completion status value is "Successful Completion" (encoded in the completion status field as "000" for this embodiment as indicated in Table 3). For example, a read request from a host bridge is routed through a switch to a completer endpoint. The Completer responds with a completion packet indicating a successful completion status and also responds with the data for the read request. The switch routes this completion packet back to the host bridge.

[0061] When a request is received and decoded by the completing device, but the completing device does not support the requested transaction and the request requires a completion, the resulting completion status is "Unsupported Request" (encoded in the

completion status field as “001” for this embodiment as indicated in Table 3). One example of an unsupported request would be a memory read request to an out-of-range address. In this case, the completer is not able to support the request and the requester is expecting a completion.

[0062] For the case where a request is received and decoded by the completing device and the completing device is unable to support the requested transaction and the requesting device is not expecting a completion, the resulting completion status is an unsupported request. Since the requesting device is not expecting a completion, the completion status is communicated to the requesting device via a message as described above in connection with Figure 10. An example of an unsupported request where the requesting device is not expecting a completion is a memory write transaction to an out-of-range address. The communication of the completion status via a message may be implemented as an optional feature.

[0063] When a completing device receives and decodes a request, but the completing device is unable to respond due to an internal error, the resulting completion status is a “Completer Abort” (encoded in the completion status field as “100” for this embodiment).

[0064] When a completing device receives a packet that violates packet formation rules, the result is a “Malformed Packet.” The completing device responds to this situation by transmitting a “Malformed Packet” error message that is routed to the requesting device. A switch that receives a malformed packet must, for this embodiment, route the packet to the upstream port if no other port can be positively identified as the intended destination port.

[0065] When a read completion has a completion status other than “Successful Completion,” no data is returned with the completion packet. The read completion with the non-successful completion status is the last completion transmitted for the request. For example, a completer may split a read request into four parts for servicing and the second completion packet results in a completer abort completion status. The final two completion packets are not transmitted. The requesting device, once it receives the completion packet with the non-successful completion status, considers the request terminated and should not expect additional completion packets corresponding to that read request.

[0066] Figure 11 is a flow diagram of one embodiment of a method for handling a completion packet that is not expected by a system agent. An “unexpected completion” occurs when an agent receives a completion that does not correspond to any outstanding requests issued by that same agent. For the example method of Figure 11, block 1110 indicates that if there is no unexpected completion then normal operation continues at block 1120. If, however, an unexpected completion is received, the unexpected completion packet is discarded at block 1130. After the packet is discarded, the above error condition may be reported to the system at block 1140. For this example embodiment, the reporting of the error may be an option that is programmable by software.

[0067] Figure 12 is a flow diagram of one embodiment of a method for a requesting device handling a completion packet with a completion status other than “Successful Completion.” Block 1210 indicates that if the completion status is “Successful Completion”, then normal operation continues at block 1220. If the completion status is other than “Successful Completion”, then at block 1230 the value of the Completer ID field is recorded. For this embodiment the Completer ID value is stored in a register. Then, at block 1240, a “Received Unsuccessful Completion” bit is set in a register in the requesting device for this embodiment. The above error condition may be reported at block 1250. The reporting of the error condition may be implemented as a programmable option. A software agent may use the Completer ID value and the “Received Unsuccessful Completion” bit to track down the source of the error condition.

[0068] Figure 13 is a flow diagram of one embodiment of a method for a completing device handling a completion packet with a completion status other than “Successful Completion.” Block 1310 indicates that if the completion status of a transmitted completion packet is “Successful Completion”, then normal operation continues at block 1320. If the completion status is other than “Successful Completion”, then at block 1330 the value of the Requester ID and Tag fields are recorded. For this embodiment, the Requester ID and Tag values are stored in one or more registers. Then, at block 1340, a “Transmitted Unsuccessful

Completion” bit is set in a register in the completing device for this embodiment. The above error condition may be reported at block 1350. The reporting of the error condition may be implemented as a programmable option. A software agent may use the Requester ID and Tag values and the “Transmitted Unsuccessful Completion” bit to track down the source of the error condition.

[0069] In the foregoing specification the invention has been described with reference to specific exemplary embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention as set forth in the appended claims. The specification and drawings are, accordingly, to be regarded in an illustrative rather than in a restrictive sense.

[0070] Reference in the specification to “an embodiment,” “one embodiment,” “some embodiments,” or “other embodiments” means that a particular feature, structure, or characteristic described in connection with the embodiments is included in at least some embodiments, but not necessarily all embodiments, of the invention. The various appearances of “an embodiment,” “one embodiment,” or “some embodiments” are not necessarily all referring to the same embodiments.

FOOTNOTES